


Forest Before Trees: Latent Superposition for Efficient Visual Reasoning

Yubo Wang^{♣*}, Juntian Zhang^{♣*}, Yichen Wu[◇], Yankai Lin[♣], Nils Lukas[♡], Yuhan Liu^{♡‡}
♡MBZUAI ♣Fudan University

♣Gaoling School of Artificial Intelligence, Renmin University of China ◇Harvard University
yubowang25@m.fudan.edu.cn {zhangjuntian, yankailin}@ruc.edu.cn
yiwu6@mg.harvard.edu {nils.lukas, yuhan.liu}@mbzuai.ac.ae

Abstract

While Chain-of-Thought empowers Large Vision-Language Models with multi-step reasoning, explicit textual rationales suffer from an information bandwidth bottleneck, where continuous visual details are discarded during discrete tokenization. Recent latent reasoning methods attempt to address this challenge, but often fall prey to premature semantic collapse due to rigid autoregressive objectives. In this paper, we propose *Laser*, a novel paradigm that reformulates visual deduction via *Dynamic Windowed Alignment Learning (DWAL)*. Instead of forcing a point-wise prediction, Laser aligns the latent state with a dynamic validity window of future semantics. This mechanism enforces a "Forest-before-Trees" cognitive hierarchy, enabling the model to maintain a probabilistic superposition of global features before narrowing down to local details. Crucially, Laser maintains interpretability via decodable trajectories while stabilizing unconstrained learning via *Self-Refined Superposition*. Extensive experiments on 6 benchmarks demonstrate that Laser achieves state-of-the-art performance among latent reasoning methods, surpassing the strong baseline Monet by 5.03% on average. Notably, it achieves these gains with extreme efficiency, reducing inference tokens by more than 97%, while demonstrating robust generalization to out-of-distribution domains. We hope this work encourages a paradigm shift from explicit next-token prediction to latent visual reasoning:  Laser

1 Introduction

Vision-Language Models (VLMs) have revolutionized visual understanding by integrating Large Language Models (LLMs) with robust visual encoders (OpenAI et al., 2024; Li et al., 2024a). While adapting Chain-of-Thought (CoT) (Wei

*Equal contribution. The order was decided by a coin flip.

†Work done during an internship at MBZUAI.

‡Corresponding author.

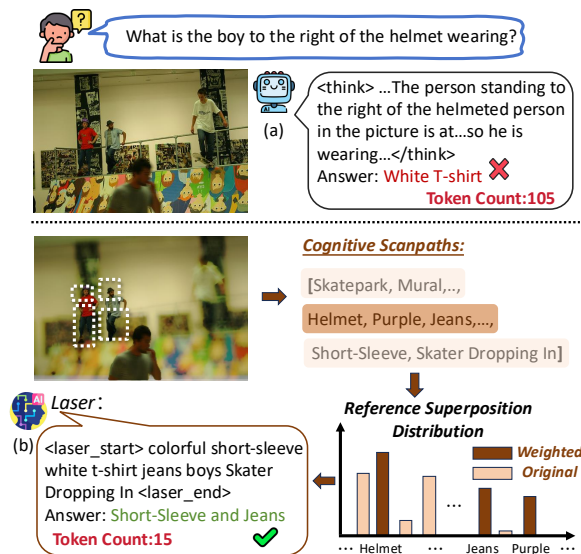


Figure 1: Laser replaces verbose textual rationales (a) with efficient latent superpositions (b). Our method achieves superior accuracy with significantly fewer tokens.

et al., 2022) has enabled multi-step reasoning, as shown in Fig 1 (a), explicit textual rationales suffer from an *information bandwidth bottleneck*, where continuous visual details are lost in discrete tokenization (Li et al., 2025). Emerging *Latent Space Reasoning* approaches (Zelikman et al., 2024; Hao et al., 2025) attempt to bypass this by reasoning within high-dimensional hidden states. However, these methods typically retain standard autoregressive objectives, forcing the latent state to strictly minimize prediction error for a specific token at every step. We argue that this strict point-wise mapping is fundamentally misaligned with visual perception. Unlike text generation, visual reasoning is hierarchical, evolving from global semantic apprehension to local feature extraction (Navon, 1977). Forcing the latent state to prematurely "collapse" into a precise object token before grasping the holistic context induces a *premature seman-*

tic collapse, creating a "tunnel vision" effect that hinders the capture of complex relationships.

To address this gap between sequential objectives and hierarchical perception, we propose **Laser (Latent Superposition for Effective Visual Reasoning)**. Laser is grounded in the insight that an effective reasoning state should not be a pointer to a single future word, but a container for a *validity domain* of future possibilities. We redefine the optimization goal via **Dynamic Windowed Alignment Learning (DWAL)**. As shown in Fig 1 (b), instead of predicting the immediate next token, Laser trains the latent state to align with a dynamic semantic window encompassing the entire remaining reasoning path. This formulation encourages the latent representation to maintain a **probabilistic superposition**, simultaneously encoding high-level global semantics while keeping specific details in a potential state. As the reasoning process unfolds, this window naturally shrinks, enforcing a progressive transition from global exploration to local precision, mimicking the general-to-specific nature of human visual processing.

Achieving this "uncollapsed" state without external supervision presents a significant optimization challenge: an unconstrained latent space risks diverging into meaningless high-entropy distributions. First, we introduce **Self-Refined Superposition**, which leverages the model's own estimation of the future window to construct a stable soft target. Second, to prevent loss of focus, we design an **Entropy-Regularized Intervention**. This mechanism acts as an implicit curriculum: it dynamically injects rigid ground-truth guidance when the model's uncertainty is high, and reverts to soft superposition when the model demonstrates a grasp of the global context.

In summary, Laser transforms visual reasoning from a rigid sequential matching task into a flexible, windowed manifold alignment problem. Our extensive experiments on complex visual reasoning benchmarks demonstrate that this approach significantly outperforms both standard CoT and baseline latent reasoning methods. By enabling VLMs to "think" in superpositions before collapsing to answers, Laser bridges the gap between the continuous nature of vision and the discrete nature of language.

Our contributions are summarized as follows:

- We propose **Laser**, a latent reasoning paradigm that reformulates visual deduction

via **Dynamic Windowed Alignment Learning (DWAL)**. This approach prevents *premature semantic collapse* by enforcing a "forest-before-trees" cognitive process within the latent space.

- We design a supervision framework combining **Self-Refined Superposition** and **Entropy-Regularized Intervention**. This establishes an implicit curriculum that stabilizes latent learning without external annotations, dynamically balancing exploration and grounding.
- We achieve a superior efficiency-performance balance: Laser obtains state-of-the-art results across 6 benchmarks while reducing inference tokens by over **97%**. Furthermore, it demonstrates robust generalization on out-of-distribution tasks, validating the efficacy of the learned visual logic.

2 Related Work

2.1 Vision-Language Models

The evolution of VLMs has rapidly advanced from static cross-modal alignment to dynamic perception. Open-source models such as LLaVA and MiniGPT-4 demonstrated that simple linear projection layers, coupled with high-quality visual instruction tuning, could achieve strong multimodal following (Liu et al., 2023; Zhu et al., 2023). To overcome perceptual limitations in fine-grained and temporal tasks, recent architectures have focused on scaling visual resolution and context length. The latest InternVL3.5, scaled vision encoders to massive parameters using dynamic tiling strategies to handle detailed imagery (Chen et al., 2024b; Wang et al., 2025c). Concurrently, Qwen2.5-VL and the advanced Qwen3-VL refined the Naive Dynamic Resolution mechanism, enabling the processing of images at arbitrary aspect ratios and extending capabilities to long-context video understanding (Bai et al., 2025b,a).

Recent research has focused extensively on enhancing VLMs' reasoning capabilities. Vision-R1 (Huang et al., 2025) and VL-Rethinker (Wang et al., 2025a) utilize Group Relative Policy Optimization, while VISC focuses on the advancement of multi-image reasoning capabilities (Zhang et al., 2025b). In parallel, architectural innovations focus on perceptual self-improvement: ViPER (Zhang et al., 2025c) establishes a self-evolutionary framework via self-critiquing cycles,

and DeepEyes (Zheng et al., 2026) integrates active perception through dynamic tool invocation. Differently from these approaches, our Laser proposes efficient latent-space reasoning.

2.2 Latent Space Reasoning

Explicit Chain-of-Thought enhances model capabilities but suffers from information loss due to discrete tokenization; consequently, LLMs like Quiet-STaR (Zelikman et al., 2024), Coconut (Hao et al., 2025), and SoftCoT (Xu et al., 2025) perform intermediate computations entirely within latent states. In VLMs, the focus shifts to anchoring latent thinking in visual evidence: Co-CoVa (Ma et al., 2026) and MCOUT (Pham and Ngo, 2025) refine representations via latent attention, whereas Mirage (Yang et al., 2025), IVTLR (Chen et al., 2025), and ILVR (Dong et al., 2025) employ interleaved decoding to stabilize reasoning. While LVR (Li et al., 2025) strengthens alignment through autoregressive reconstruction, it risks representation collapse; alternatively, “visual scratchpads” like Latent Sketchpad (Zhang et al., 2025a) and SkiLa (Tong et al., 2025) preserve interpretability via reconstructable latent sketches.

Beyond supervised anchoring, other methods directly optimize latent reasoning trajectories (Monet (Wang et al., 2025b), LaCoT (Sun et al., 2025), DMLR (Liu et al., 2025)). Mull-Tokens (Ray et al., 2025) offers a modality-agnostic workspace, while Titans (Behrouz et al., 2024) targets long-range dependencies. Distinct from these, our **Laser** method introduces a *Dynamic Windowed Alignment mechanism*. By encoding global visual semantics into a compact superposition state, it achieves a superior balance between latent reasoning and efficiency.

3 Methodology

In this section, we propose **Laser**, an efficient visual reasoning method operating in the latent space. Section 3.1 provides a formal definition of the problem. The key to realizing Laser lies in data acquisition and the design of the training methodology, which are elaborated in Section 3.2 and Section 3.3 respectively.

3.1 Problem Formulation

We formulate the visual reasoning task as a two-stage conditional generation process: *Latent Visual Reasoning* followed by *Explicit Answer Generation*. Given an input image I and a textual query

Q , the model \mathcal{M}_θ aims to synthesize a chain of visual concepts $\mathcal{C} = \{c_1, c_2, \dots, c_T\}$ that acts as an intermediate reasoning path, before producing the final response $\mathcal{A} = \{a_1, a_2, \dots, a_M\}$.

The core of our formulation focuses on the **Latent Reasoning Trajectory**. The model processes the multimodal context to generate a sequence of high-dimensional hidden states $\mathcal{H} = \{h_1, h_2, \dots, h_T\}$, where $h_t = \mathcal{M}_\theta(I, Q, c_{<t})$. This latent state is projected onto the vocabulary space via a linear head to obtain the logits $z_t = \mathbf{W}_u h_t$. These logits define the probability distribution over the next token via the Softmax function:

$$P_\theta(c_{t+1} | I, Q, c_{<t}) = \text{Softmax}(z_t). \quad (1)$$

In standard autoregressive frameworks, the optimization objective strictly forces each latent state h_t to minimize the negative log-likelihood of the immediate next token c_{t+1} . However, this rigid, local constraint compels the latent representation to collapse prematurely into a single semantic point. In this work, we redefine the optimization goal: rather than solely predicting the next token, the latent state h_t is tasked with aligning with a *dynamic validity domain* of future visual semantics. Thus, the problem transforms from minimizing point-wise prediction error to maximizing the **Windowed Semantic Alignment** between the latent trajectory \mathcal{H} and the reasoning chain \mathcal{C} , laying the theoretical foundation for the Dynamic Windowed Alignment Learning (DWAL) detailed in Sec. 3.3.

3.2 Synthesizing Cognitive Scanpaths

For the Laser method, we require a dataset that captures the intermediate visual reasoning process without relying on expensive human annotations. Unlike previous approaches such as Visual CoT (Shao et al., 2024), which anchor reasoning chains with explicit bounding boxes, our aim is to bridge perception and language through *implicit* latent alignment. Consequently, we construct a scalable annotation pipeline that operates under a weakly supervised setting: relying solely on synthesized semantic sequences while deliberately excluding explicit Region-of-Interest (ROI) supervision.

We leverage GPT-4o as a “Visual Cognitive Engine” to sequentially synthesize reasoning paths composed of discrete semantic tokens. Crucially, our prompt design is grounded in the *Global Precedence Hypothesis* (Navon, 1977), which posits

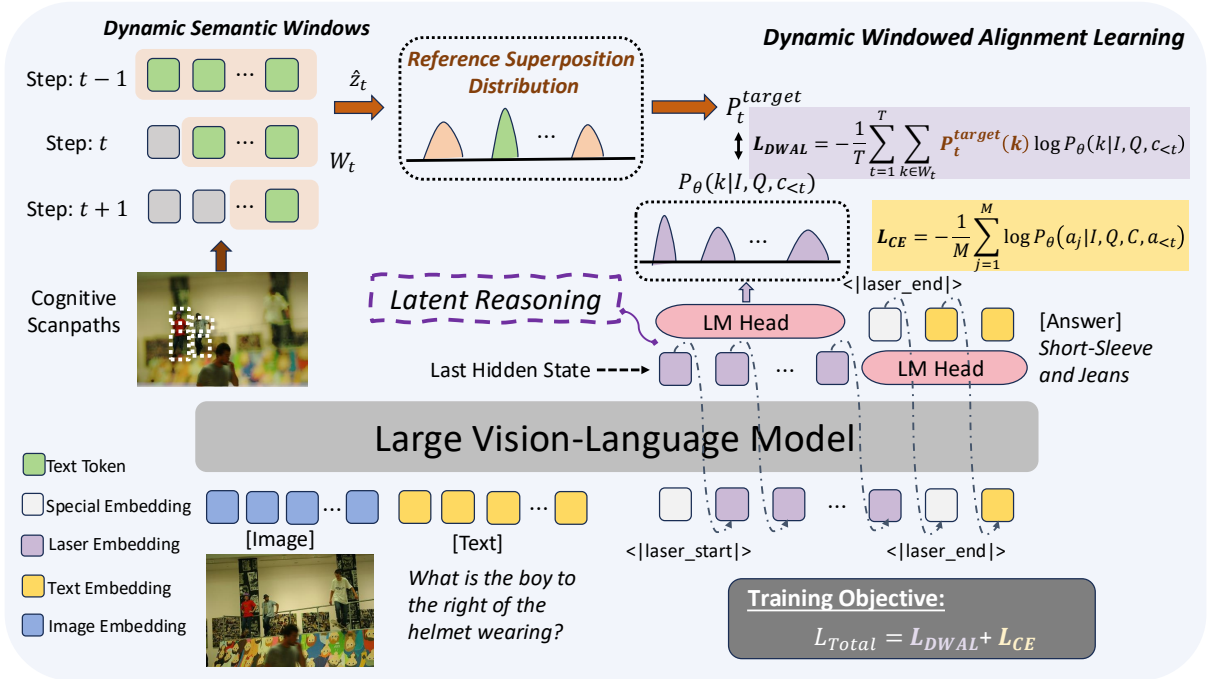


Figure 2: **Overview of the Laser.** Laser employs DWAL. At each step t , a dynamic validity window \mathcal{W}_t is defined over future semantic tokens to construct a **Reference Superposition Distribution**. The latent state is then optimized to align with this distribution via \mathcal{L}_{DWAL} . The final answer is generated explicitly after the reasoning using \mathcal{L}_{CE} .

that human perception inherently prioritizes holistic structures (“forest”) before processing detailed components (“trees”). Guided by this principle, we enforce a strict *Global-to-Local Scanning Logic*: the synthesized sequences must initiate with a global anchor, progressively narrow focus to relevant objects, and culminate in the critical visual evidence required to answer the query. This structure ensures that the data represent valid deductive trajectories rather than static descriptions. After applying a rigorous filtering protocol, we obtain a final dataset named *ScanPath*, consisting of 270k high-quality samples. More details on prompt engineering and filtering statistics are provided in the Appendix F and H.

3.3 Dynamic Windowed Alignment Learning

At the core of **Laser** is the **Dynamic Windowed Alignment Learning (DWAL)**. Standard autoregressive objectives typically enforce a rigid, point-wise collapse to a single ground-truth token at each step, which we argue is suboptimal for early-stage visual reasoning where global context is paramount. In contrast, DWAL formulates reasoning as a Windowed Probabilistic Alignment problem. By defining a dynamic validity domain of future semantics, this approach allows the latent state to encode a superposition of potential visual concepts, mimick-

ing the general-to-specific nature of human visual processing.

3.3.1 Dynamic Semantic Windows

Let $\mathcal{H} = \{h_1, h_2, \dots, h_T\}$ denote the sequence of latent hidden states generated by the model during the reasoning phase, and let $\mathcal{C} = \{c_1, c_2, \dots, c_T\}$ denote the corresponding sequence of text tokens representing visual concepts, which are annotated during data construction.

Previous approaches (Li et al., 2025) typically enforce a *Strict Point-wise Mapping*, minimizing the divergence $\mathcal{D}(h_t, c_t)$ between the latent state h_t and the specific ground-truth visual concept c_t at every timestep t . To overcome the limitations of this premature semantic collapse, we introduce a **Dynamic Semantic Window** \mathcal{W}_t for each reasoning step t . This window defines the “valid semantic field” that the current state should encompass:

$$\mathcal{W}_t = \{c_k \mid t \leq k \leq T\}. \quad (2)$$

The objective is to encourage the latent representation h_t to cover the full spectrum of the valid window \mathcal{W}_t , rather than peaking solely at the immediate next token c_{t+1} . As t increases, the window \mathcal{W}_t naturally shrinks ($|\mathcal{W}_t| \rightarrow 1$), enforcing a progressive transition from global semantic superposition to local precision.

Note that to ensure the reasoning process terminates effectively, the dedicated special token `<laser_end>` is explicitly excluded from the validity window W_t for all steps $t \in [1, T]$. Instead, `<laser_end>` serves as a deterministic target only after the completion of the final reasoning step T , signaling the phase transition from implicit reasoning to explicit answer generation.

3.3.2 Learning via Latent Superposition

To supervise the model within the dynamic window W_t without relying on external soft labels, we employ a **Self-Refined Superposition** mechanism. This approach leverages the model’s own estimation of the valid semantic manifold to construct a stable soft target. Specifically, we extract the logits corresponding to the tokens in W_t and apply a *stop-gradient* operation to prevent unstable self-reinforcement loops. Let $\hat{z}_t^{(k)} = \text{StopGrad}(z_t^{(k)})$ denote the detached logit for token $c_k \in W_t$. We define a **reference superposition distribution** Q_t via a temperature-scaled Softmax:

$$Q_t(k) = \frac{\exp(\hat{z}_t^{(k)}/\tau)}{\sum_{j \in W_t} \exp(\hat{z}_t^{(j)}/\tau)}, \quad (3)$$

where τ is a hyperparameter controlling the sharpness of the distribution. This formulation encourages the hidden state h_t to maintain a probabilistic superposition of future visual concepts.

However, relying solely on soft targets can lead to optimization divergence, where the model converges to a high-entropy uniform distribution lacking semantic focus. To mitigate this, we introduce an **Entropy-Regularized Intervention**. We first compute the normalized entropy of the reference distribution to gauge the model’s uncertainty:

$$H(Q_t) = -\frac{1}{\log |W_t|} \sum_{k \in W_t} Q_t(k) \log Q_t(k). \quad (4)$$

We then construct a hybrid target P_t^{target} that dynamically switches between the soft superposition and a rigid next-token alignment based on this uncertainty:

$$P_t^{\text{target}} = \begin{cases} \alpha \cdot \mathbf{y}_{\text{hard}} + (1 - \alpha) \cdot Q_t, & \text{if } H(Q_t) > \eta \\ Q_t, & \text{otherwise} \end{cases} \quad (5)$$

where \mathbf{y}_{hard} is the one-hot vector for the immediate next token c_{t+1} , η is a predefined entropy threshold, and $\alpha \in [0, 1]$ controls the intensity of the

hard intervention. This mechanism creates an *implicit curriculum*: it enforces precise grounding when the model exhibits high uncertainty (high entropy), while enabling superposition-based reasoning when the model has grasped the global context.

3.3.3 Optimization Objective

The total optimization objective unifies the latent reasoning process and the explicit answer generation. For the reasoning chain, the DWAL Loss minimizes the cross-entropy between the hybrid target and the model’s prediction, effectively aligning the latent trajectory with the dynamic semantic windows:

$$\mathcal{L}_{\text{DWAL}} = -\frac{1}{T} \sum_{t=1}^T \sum_{k \in W_t} P_t^{\text{target}}(k) \log P_\theta(k | I, Q, c_{<t}). \quad (6)$$

Subsequently, for the answer generation phase, the model produces the final response tokens $\mathcal{A} = \{a_j\}_{j=1}^M$ based on the evolved visual understanding. We adopt the standard Cross-Entropy (CE) loss, conditioning on the image I , the original query Q , and the completed visual chain \mathcal{C} :

$$\mathcal{L}_{\text{CE}} = -\frac{1}{M} \sum_{j=1}^M \log P_\theta(a_j | I, Q, \mathcal{C}, a_{<j}). \quad (7)$$

The final training objective is a weighted summation of these two components:

$$\mathcal{L}_{\text{Total}} = \mathcal{L}_{\text{DWAL}} + \mathcal{L}_{\text{CE}}. \quad (8)$$

By minimizing $\mathcal{L}_{\text{Total}}$, **Laser** effectively balances the exploration of global visual semantics during the reasoning phase with the exploitation of precise local semantics for answer generation.

4 Experiments

We first demonstrate the superiority of **Laser** through extensive evaluations on 6 diverse benchmarks against state-of-the-art models. We then conduct further investigations into our method through multi-faceted studies and in-depth analysis.

4.1 Experimental Setup

We instantiate **Laser** using Qwen2.5-VL-7B-Instruct as the backbone. To preserve the pre-trained visual representations and ensure training efficiency, we freeze the vision tower and the modality merger, exclusively optimizing the LLM

Model	MMVP	BLINK	SEEDBENCH2PLUS	MMStar	Hallusion Bench	HRBENCH	Overall
<i>Zero-Shot VLMs</i>							
GPT-4o (OpenAI et al., 2024)	68.70	68.00	72.00	64.70	–	–	–
Qwen2.5-VL-7B (Bai et al., 2025b)	65.67	53.60	65.31	59.70	56.57	68.25	61.52
LLaVA-OneVision (Li et al., 2024a)	74.00	49.34	61.22	59.13	51.10	63.00	59.63
InternVL3.5-8B (Wang et al., 2025c)	57.67	54.81	69.78	53.33	56.15	59.38	58.52
<i>Tool-use & RL Enhanced Reasoning</i>							
PAPO (Wang et al., 2026)	68.67	52.66	54.11	45.80	57.52	68.12	57.81
Vision-R1 (Huang et al., 2025)	72.67	52.71	68.95	62.67	63.83	75.12	65.99
VL-Rethinker (Wang et al., 2025a)	72.67	55.55	70.27	63.20	71.08	63.50	66.05
DeepEyes (Zheng et al., 2026)	70.00	51.08	69.08	58.73	62.57	69.12	63.43
<i>Latent Reasoning</i>							
Monet (Wang et al., 2025b)	68.00	50.71	<u>65.88</u>	60.33	56.36	68.00	61.55
LVR (Li et al., 2025)	64.00	<u>53.60</u>	47.39	57.93	<u>65.19</u>	53.62	56.96
Laser (Ours)	72.00	56.92	70.05	<u>60.27</u>	67.72	72.50	66.58
$\Delta \uparrow$	4.00	6.21	4.17	-0.33	11.36	4.50	5.03

Table 1: Main results comparing **Laser** with baselines across three paradigms: Zero-Shot VLMs, Tool-use & RL, and Latent Reasoning. The best results among *latent reasoning* methods are highlighted in **bold**, and the second best are underlined. Δ denotes the absolute performance gain over the strongest latent baseline, Monet (Wang et al., 2025b).

parameters. Regarding our specific **Laser** configuration, we set the temperature $\tau = 1.0$ to modulate the softness of the reference distribution, and the entropy threshold $\eta = 0.6$ to control the intervention mechanism. Comprehensive hyperparameters are detailed in Appendix A.

4.2 Baselines

We evaluate Laser against a comprehensive set of state-of-the-art baselines across three paradigms: (1) **Zero-shot VLMs** (GPT-4o (OpenAI et al., 2024), LLaVA-OneVision (Li et al., 2024a), InternVL3.5-8B (Wang et al., 2025c), Qwen2.5-VL-7B (Bai et al., 2025b)), (2) Explicit Visual Interaction methods, including **tool-augmented reasoning** (DeepEyes (Zheng et al., 2026)) and **RL-enhanced VLM reasoning** (Vision-R1 (Huang et al., 2025), PAPO (Wang et al., 2026), VL-Rethinker (Wang et al., 2025a), and (3) **Latent VLM Reasoning** approaches (LVR (Li et al., 2025), Monet (Wang et al., 2025b)). Please refer to Appendix B for more details.

4.3 Benchmarks

We evaluate Laser on six comprehensive benchmarks covering diverse scenarios. For *visual perception*, we use **BLINK** (Fu et al., 2024) to stress image-dependent perception such as depth and spatial cues, and **MMVP** (Tong et al., 2024) to probe CLIP-blind visual patterns. For *visual reasoning*, we adopt **MMStar** (Chen et al., 2024a), which evaluates fine-grained reasoning axes while mini-

mizing shortcut leakage. To probe *high-resolution understanding*, we include **HRBench** (Wang et al., 2025d) for ultra-high-resolution (4K only) visual perception. Moreover, we assess *trustworthiness and text-rich understanding* using **HallusionBench** (Guan et al., 2024) to diagnose visual illusion and language hallucinations (measured by Question Accuracy, Q-Acc), and **SEED-Bench-2-Plus** (Li et al., 2024b) to evaluate comprehensive understanding of text-intensive visuals such as charts, maps, and web pages. Each benchmark are detailed in Appendix C.

4.4 Results

The comparative results across six benchmarks are presented in Table 1. Laser establishes a new state-of-the-art among latent reasoning methods and demonstrates superior capabilities even against computationally intensive explicit reasoning paradigms.

As our primary focus, Laser significantly outperforms existing latent reasoning baselines. Compared to the previous best method, Monet, Laser achieves a remarkable +5.03% gain in the overall score. Notably, we observe the most substantial improvements on HallusionBench (+11.36%) and BLINK (+6.21%). We attribute these gains to the proposed Dynamic Windowed Alignment. By maintaining a semantic superposition rather than collapsing to a rigid token prematurely, Laser effectively mitigates the hallucination issues common in point-wise latent methods and captures fine-grained

visual details. In contrast, LVR, which enforces strict next-token reconstruction, lags significantly behind (-9.62%), highlighting the necessity of our flexible windowed strategy.

Beyond the latent domain, Laser compares favorably against heavyweight paradigms. Despite operating purely in the compact latent space without external tools or iterative reinforcement learning search, Laser surpasses both the leading RL-based method, Vision-R1, and the tool-augmented VL-Rethinker. This suggests that optimizing the internal cognitive trajectory via superposition is a more efficient path to enhanced reasoning than externalizing the process into lengthy textual chains. Laser consistently outperforms its backbone model, Qwen2.5-VL-7B, across all evaluated benchmarks. Specifically, on the MMVP benchmark, which tests CLIP-blind patterns, Laser improves upon the baseline by +6.33%. This demonstrates that our method is effectively unlocks latent visual discrimination capabilities that are otherwise dormant in standard supervision settings.

5 Discussions

In this section, we address five key research questions through systematic investigation, providing deeper insights into Laser’s performance and behavior.

RQ1: How efficient is Laser compared to baselines?

Model	Blink ($N=1901$)		HrBench ($N=800$)	
	Avg Tokens	Δ	Avg Tokens	Δ
Qwen2.5-VL-7B	223.5	–	55.9	–
VL-Rethinker	207.0	↓-7.4%	143.8	↑+157.2%
Monet-7B	118.3	↓-47.1%	86.8	↑+55.3%
LVR	8.0	↓-96.4%	8.0	↓-85.7%
Laser	6.0	↓-97.3%	5.7	↓-89.7%

Table 2: Efficiency comparison on Blink and HrBench. Our **Laser** achieves a significant reduction in token usage. The light yellow background indicates *Explicit Reasoning* methods, while the light blue background represents *Latent Reasoning* methods. N denotes the total number of samples, Δ represents the relative change compared to the Qwen2.5-VL-7B, and ↓ indicates efficiency improvement.

Beyond raw accuracy, the practical deployment of VLMs is often constrained by inference latency and computational cost. We analyze the inference efficiency of **Laser** compared to both standard baselines and recent latent reasoning methods on the BLINK and HRBench. The results are summarized

in Table 2.

Existing explicit reasoning approaches, such as VL-Rethinker, typically rely on generating lengthy textual rationales to bridge the gap between perception and reasoning. As shown in Table 2, this imposes a severe computational burden. For instance, on HRBench, VL-Rethinker increases token consumption by **157.2%** compared to the base model, directly leading to high inference latency. Interestingly, while Monet is designed as a latent reasoning framework, it still incurs a substantial computational overhead. On BLINK, Monet uses **118.3 tokens** on average. Although this is a reduction compared to the base model, it remains orders of magnitude heavier than other latent approaches, suggesting that its “visual thoughts” still occupy a dense latent sequence.

In contrast, **Laser** achieves exceptional efficiency by shifting the reasoning process from the discrete token space to the continuous latent space, significantly outperforming both explicit and latent baselines. On the BLINK benchmark, Laser reduces the average token count to merely **6.0 tokens**, a reduction of **97.3%**. This makes Laser substantially more efficient than Monet (118.3 tokens) and even surpasses the comparable latent method LVR (8.0 tokens).

Crucially, as noted in Table 1, this efficiency does not compromise performance. While LVR suffers from semantic degradation due to its strict reconstruction objective, Laser’s superposition mechanism allows it to encode richer, non-collapsed semantic information within the same compact latent budget. This confirms that Laser achieves a superior trade-off between efficiency and accuracy: it delivers the reasoning depth of explicit Chain-of-Thought models while maintaining the inference speed of direct-answer models. By condensing reasoning into a compact superposition state, Laser eliminates the need for generating hundreds of intermediate tokens, offering near-instantaneous inference suitable for real-time applications.

RQ2: What is Laser’s impact across variant tasks?

We further conducted an in-depth experimental analysis of Laser across a range of task domains covering perception, understanding, and reasoning. As shown in Figure 3, Laser exhibits a distinct performance profile: it achieves dominant superiority in 11 out of 14 tasks, with the only exceptions being *Object Localization*, *Jigsaw*, and *Functional Correspondence*. This polarization reveals a fun-

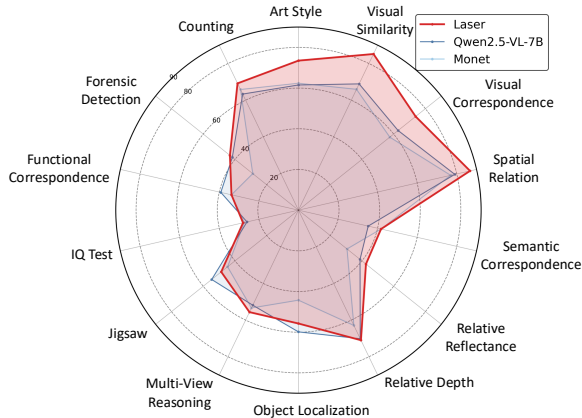


Figure 3: Fine-grained comparison across 14 distinct categories. Laser outperforms Qwen2.5-VL-7B and Monet in 11 tasks, highlighting superior high-level semantic and spatial reasoning.

damental insight into the "Forest Before Trees" cognitive mechanism. Laser excels in tasks demanding high-level discrimination, such as *Visual Similarity* and *Spatial Relation*, by avoiding the premature information loss inherent in rigid tokenization. By preserving a "probabilistic haze" of visual nuances, the model effectively soft-matches ambiguous patterns and grasps relative 3D geometry. This confirms that our Dynamic Semantic Window successfully captures the holistic "Forest" structure, prioritizing scene-level contextual coherence over isolated object identification.

Conversely, Laser slightly underperforms in *Object Localization* and *Jigsaw*. This creates an intriguing contrast: while Laser is superior at understanding spatial relations, it is less precise at absolute pixel-level grounding. We attribute this to our weakly-supervised design choice. Unlike methods that rely on explicit bounding box regression (Region-of-Interest supervision), Laser learns grounding implicitly via latent alignment. It prioritizes the semantic flow of reasoning over rigid pixel reconstruction. Tasks like *Jigsaw* require exact low-level feature matching rather than abstract semantic reasoning. Laser’s “Forest-first” strategy naturally favors holistic scene understanding, leading the model to abstract away fine-grained pixel details. While this results in a minor trade-off in absolute localization precision, it yields the significant robustness observed in complex reasoning scenarios, mimicking human cognition, which is often semantically precise but metrically approximate.

RQ3: Does Laser influence general vision-

language abilities?

Model	Multi-View Reasoning	Relative Depth	Geometry	Math	Web	Chart
Qwen2.5-VL-7B	51.88	70.16	53.24	66.00	75.45	61.98
Laser (Ours)	55.64	70.97	53.24	67.20	83.48	67.16

Table 3: We compare Laser against Qwen2.5-VL-7B across three dimensions: *Spatial Perception* (Multi-View, Depth, Geometry), *Visual Logic* (Math), and *Structural Understanding* (Web, Charts). The results confirm that Laser successfully transfers reasoning patterns to unseen domains while preserving general spatial capabilities

Addressing the pervasive risk of catastrophic forgetting, we evaluate whether Laser maintains its general-purpose visual foundation. The results on out-of-distribution domains¹ confirm that our method avoids performance degradation and effectively preserves general skills. Specifically, Laser achieves substantial gains of 8.03% on Web and 5.18% on Chart tasks while maintaining a steady score of 53.24 on Geometry, suggesting that our "Forest-before-Trees" hierarchy reinforces the grasp of global structures. Furthermore, the learned deductive patterns successfully transfer to unseen logic domains, evidenced by improvements of 1.20% in Math and 0.81% in Relative Depth. As shown in Table 3, Laser boosts specialized reasoning without compromising the model’s robust general capabilities.

RQ4: Is Laser’s reasoning interpretable?



Q: Where is the seated person visible in the image?

- A: On the fence.
- B: In the field
- C: Outside the field. ✓
- D: Behind the bench

Step	Token				
0	Se	Spect	Seats	Crowd	Audience
1	Field	Fenc	Behind	Out	
2	Outside	C	< laser_end >		

Figure 4: Visualization of the latent cognitive trajectory. The decoded tokens reveal a structured multi-hop reasoning path, evolving from entity localization (Step 0: Seats) to spatial analysis (Step 1: Fence) and final deduction.

A pivotal advantage of Laser is its **inherent interpretability**, derived from the rigorous alignment

¹Task sources: Multi-View and Relative Depth (**Blink**), Geometry (**Geo**), Math (**MMStar**), and Web/Chart (**SEED-Bench-2+**).

between the visual projector and the LLM’s semantic space. Unlike the opaque continuous vectors found in standard latent reasoning models, Laser’s hidden states can be directly projected onto the vocabulary via the frozen LM head. This allows us to inspect the top- k tokens at each intermediate step, effectively visualizing the model’s “cognitive trajectory.” We demonstrate this capability using a representative case from MMStar, where the model must locate “the seated person” in a baseball scene dominated by a salient pitcher in the foreground. **Initially (Step 0)**, the decoded tokens, led by “Se-” (suggesting *Seats*), “Spect-” (*Spectators*), and “Crowd”, reveal that the model successfully overcomes visual saliency bias, shifting its attention to the background audience. **Subsequently (Step 1)**, the latent state evolves to encode spatial constraints, with tokens such as “Fence,” “Behind,” and “Out” emerging to define the boundary between the spectators and the field. **Finally (Step 2)**, this spatial reasoning converges into a semantic decision, as the probability mass shifts to “Outside” and the correct option label “C.” This trajectory confirms that Laser performs *explicit-like* multi-hop reasoning, progressing from entity localization to spatial analysis and final deduction, entirely within the compact latent space.

RQ5: How critical is each component?




Figure 5: Ablation study. We contrast the full Laser model with variants lacking the DWAL (w/o DWAL) and the dynamic windows (w/o Windows). The consistent performance gap across six benchmarks highlights the necessity of the proposed Dynamic Windowed Alignment Learning for effective visual reasoning.

To investigate the mechanisms behind the performance gains of Laser, we conduct an ablation study isolating the superposition objective and the windowing strategy. As illustrated in Figure 5, the results reveal that these components serve distinct roles across different task types. First, removing the DWAL objective (reverting to standard next-token prediction) leads to a significant drop in fine-

grained perception benchmarks. This empirically validates that probabilistic superposition is crucial for preventing the “premature semantic collapse” often observed in standard autoregressive models. Second, using a fixed validity window (“w/o Windows”) primarily impairs performance on complex reasoning tasks, with less impact on pure perception. This confirms that the Dynamic Window strategy is essential for enforcing the “Forest-before-Trees” hierarchy. By progressively shrinking the semantic scope, it ensures the model captures global context before focusing on local details.

6 Conclusion

We introduce *Laser*, a method that transcends discrete Chain-of-Thought via continuous latent superposition. By reformulating visual deduction with Dynamic Windowed Alignment Learning, we enforce a “Forest-before-Trees” hierarchy that prevents premature semantic collapse. Laser achieves state-of-the-art performance among latent reasoning methods with superior robustness and a 97% reduction in inference overhead. Notably, it stands as the first interpretable latent reasoning approach, suggesting that emancipating reasoning from rigid tokenization fosters native multimodal intelligence.

Limitations

Despite the promising results, we acknowledge several limitations inherent to our current study. A primary constraint lies in the scope of data and computational resources. The efficacy of our approach depends on high-quality reasoning annotations, yet the scarcity of diverse, large-scale multimodal datasets limits our current task coverage. Additionally, to maintain a balance between training efficiency and computational accessibility, we restricted our study to standard-resolution static images. Consequently, the applicability of Laser to computationally demanding scenarios, such as high-resolution document analysis or temporal video reasoning, remains to be explored.

Furthermore, the evaluation of implicit reasoning faces methodological challenges. Current benchmarks are predominantly outcome-oriented, lacking standardized metrics to assess the quality of the intermediate reasoning process itself. Although Laser improves structural transparency by aligning internal states with human perceptual patterns, its reasoning trajectory remains encoded in high-dimensional vectors. This distinct modality pre-

cludes the application of established textual evaluation metrics, making it difficult to quantitatively verify the correctness of the intermediate steps, which is a significant hurdle for the rigorous assessment of latent reasoning paradigms.

Acknowledgments

This work was supported by funding from Mohamed bin Zayed University of Artificial Intelligence (MBZUAI).

References

- Shuai Bai, Yuxuan Cai, Ruizhe Chen, Keqin Chen, Xionghui Chen, Zesen Cheng, Lianghao Deng, Wei Ding, Chang Gao, Chunjiang Ge, Wenbin Ge, Zhifang Guo, Qidong Huang, Jie Huang, Fei Huang, Binyuan Hui, Shutong Jiang, Zhaohai Li, Mingsheng Li, and 45 others. 2025a. [Qwen3-vl technical report](#). *Preprint*, arXiv:2511.21631.
- Shuai Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, Sibao Song, Kai Dang, Peng Wang, Shijie Wang, Jun Tang, Humen Zhong, Yuanzhi Zhu, Mingkun Yang, Zhaohai Li, Jianqiang Wan, Pengfei Wang, Wei Ding, Zheren Fu, Yiheng Xu, and 8 others. 2025b. [Qwen2.5-vl technical report](#). *Preprint*, arXiv:2502.13923.
- Ali Behrouz, Peilin Zhong, and Vahab Mirrokni. 2024. [Titans: Learning to memorize at test time](#). *Preprint*, arXiv:2501.00663.
- Chao Chen, Zhixin Ma, Yongqi Li, Yupeng Hu, Yinwei Wei, Wenjie Li, and Liqiang Nie. 2025. [Reasoning in the dark: Interleaved vision-text reasoning in latent space](#). *Preprint*, arXiv:2510.12603.
- Lin Chen, Jinsong Li, Xiaoyi Dong, Pan Zhang, Yuhang Zang, Zehui Chen, Haodong Duan, Jiaqi Wang, Yu Qiao, Dahua Lin, and Feng Zhao. 2024a. [Are we on the right way for evaluating large vision-language models?](#) In *Advances in Neural Information Processing Systems*, volume 37, pages 27056–27087. Curran Associates, Inc.
- Zhe Chen, Jiannan Wu, Wenhai Wang, Weijie Su, Guo Chen, Sen Xing, Muyan Zhong, Qinglong Zhang, Xizhou Zhu, Lewei Lu, Bin Li, Ping Luo, Tong Lu, Yu Qiao, and Jifeng Dai. 2024b. [Internvl: Scaling up vision foundation models and aligning for generic visual-linguistic tasks](#). In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 24185–24198.
- Shuai Dong, Siyuan Wang, Xingyu Liu, and Zhongyu Wei. 2025. [Interleaved latent visual reasoning with selective perceptual modeling](#). *Preprint*, arXiv:2512.05665.
- Xingyu Fu, Yushi Hu, Bangzheng Li, Yu Feng, Haoyu Wang, Xudong Lin, Dan Roth, Noah A Smith, Wei-Chiu Ma, and Ranjay Krishna. 2024. [Blink: Multimodal large language models can see but not perceive](#). In *European Conference on Computer Vision*, pages 148–166. Springer.
- Tianrui Guan, Fuxiao Liu, Xiyang Wu, Ruiqi Xian, Zongxia Li, Xiaoyu Liu, Xijun Wang, Lichang Chen, Furong Huang, Yaser Yacoob, Dinesh Manocha, and Tianyi Zhou. 2024. [Hallusionbench: an advanced diagnostic suite for entangled language hallucination and visual illusion in large vision-language models](#). In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14375–14385.
- Shibo Hao, Sainbayar Sukhbaatar, DiJia Su, Xian Li, Zhiting Hu, Jason Weston, and Yuandong Tian. 2025. [Training large language models to reason in a continuous latent space](#). *Preprint*, arXiv:2412.06769.
- Wenxuan Huang, Bohan Jia, Zijie Zhai, Shaosheng Cao, Zheyu Ye, Fei Zhao, Zhe Xu, Yao Hu, and Shaohui Lin. 2025. [Vision-r1: Incentivizing reasoning capability in multimodal large language models](#). *arXiv preprint arXiv:2503.06749*.
- Bangzheng Li, Ximeng Sun, Jiang Liu, Ze Wang, Jialian Wu, Xiaodong Yu, Hao Chen, Emad Barsoum, Muhao Chen, and Zicheng Liu. 2025. [Latent visual reasoning](#). *arXiv preprint arXiv:2509.24251*.
- Bo Li, Yuanhan Zhang, Dong Guo, Renrui Zhang, Feng Li, Hao Zhang, Kaichen Zhang, Peiyuan Zhang, Yanwei Li, Ziwei Liu, and Chunyuan Li. 2024a. [Llava-onevision: Easy visual task transfer](#). *Preprint*, arXiv:2408.03326.
- Bohao Li, Yuying Ge, Yixiao Ge, Guangzhi Wang, Rui Wang, Ruimao Zhang, and Ying Shan. 2024b. [SEED-Bench-2-Plus: Benchmarking multimodal large language models with text-rich visual comprehension](#). *arXiv preprint arXiv:2404.16790*.
- Chengzhi Liu, Yuzhe Yang, Yue Fan, Qingyue Wei, Sheng Liu, and Xin Eric Wang. 2025. [Reasoning within the mind: Dynamic multimodal interleaving in latent space](#). *Preprint*, arXiv:2512.12623.
- Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. 2023. [Visual instruction tuning](#). In *Advances in Neural Information Processing Systems*, volume 36, pages 34892–34916. Curran Associates, Inc.
- Jizheng Ma, Xiaofei Zhou, Geyuan Zhang, Yanlong Song, and Han Yan. 2026. [Multimodal reasoning via latent refocusing](#). *Preprint*, arXiv:2511.02360.
- David Navon. 1977. [Forest before trees: The precedence of global features in visual perception](#). *Cognitive psychology*, 9(3):353–383.
- OpenAI, Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, Red Avila, Igor Babuschkin,

- Suchir Balaji, Valerie Balcom, Paul Baltescu, Haiming Bao, Mohammad Bavarian, Jeff Belgum, and 262 others. 2024. [Gpt-4 technical report](#). *Preprint*, arXiv:2303.08774.
- Tan-Hanh Pham and Chris Ngo. 2025. [Multi-modal chain of continuous thought for latent-space reasoning in vision-language models](#). *Preprint*, arXiv:2508.12587.
- Arijit Ray, Ahmed Abdelkader, Chengzhi Mao, Bryan A. Plummer, Kate Saenko, Ranjay Krishna, Leonidas Guibas, and Wen-Sheng Chu. 2025. [Mull-tokens: Modality-agnostic latent thinking](#). *Preprint*, arXiv:2512.10941.
- Hao Shao, Shengju Qian, Han Xiao, Guanglu Song, Zhuofan Zong, Letian Wang, Yu Liu, and Hongsheng Li. 2024. [Visual cot: Advancing multi-modal language models with a comprehensive dataset and benchmark for chain-of-thought reasoning](#). In *Advances in Neural Information Processing Systems*, volume 37, pages 8612–8642. Curran Associates, Inc.
- Guohao Sun, Hang Hua, Jian Wang, Jiebo Luo, Soheil Dianat, Majid Rabbani, Raghuvver Rao, and Zhiqiang Tao. 2025. [Latent chain-of-thought for visual reasoning](#). *Preprint*, arXiv:2510.23925.
- Jintao Tong, Jiaqi Gu, Yujing Lou, Lubin Fan, Yixiong Zou, Yue Wu, Jieping Ye, and Ruixuan Li. 2025. [Sketch-in-latents: Eliciting unified reasoning in mllms](#). *Preprint*, arXiv:2512.16584.
- Shengbang Tong, Zhuang Liu, Yuexiang Zhai, Yi Ma, Yann LeCun, and Saining Xie. 2024. [Eyes wide shut? exploring the visual shortcomings of multimodal llms](#). In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 9568–9578.
- Haozhe Wang, Chao Qu, Zuming Huang, Wei Chu, Fangzhen Lin, and Wenhua Chen. 2025a. [Vl-rethinker: Incentivizing self-reflection of vision-language models with reinforcement learning](#). *Preprint*, arXiv:2504.08837.
- Qixun Wang, Yang Shi, Yifei Wang, Yuanxing Zhang, Pengfei Wan, Kun Gai, Xianghua Ying, and Yisen Wang. 2025b. [Monet: Reasoning in latent visual space beyond images and language](#). *Preprint*, arXiv:2511.21395.
- Weiyun Wang, Zhangwei Gao, Lixin Gu, Hengjun Pu, Long Cui, Xingguang Wei, Zhaoyang Liu, Linglin Jing, Shenglong Ye, Jie Shao, Zhaokai Wang, Zhe Chen, Hongjie Zhang, Ganlin Yang, Haomin Wang, Qi Wei, Jinhui Yin, Wenhao Li, Erfei Cui, and 56 others. 2025c. [InternV3.5: Advancing open-source multimodal models in versatility, reasoning, and efficiency](#). *Preprint*, arXiv:2508.18265.
- Wenbin Wang, Liang Ding, Minyan Zeng, Xiabin Zhou, Li Shen, Yong Luo, Wei Yu, and Dacheng Tao. 2025d. [Divide, conquer and combine: A training-free framework for high-resolution image perception in multimodal large language models](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 39(8):7907–7915.
- Zhenhailong Wang, Xuehang Guo, Sofia Stoica, Haiyang Xu, Hongru Wang, Hyeonjeong Ha, Xiusi Chen, Yangyi Chen, Ming Yan, Fei Huang, and Heng Ji. 2026. [Perception-aware policy optimization for multimodal reasoning](#). *Preprint*, arXiv:2507.06448.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, brian ichter, Fei Xia, Ed Chi, Quoc V Le, and Denny Zhou. 2022. [Chain-of-thought prompting elicits reasoning in large language models](#). In *Advances in Neural Information Processing Systems*, volume 35, pages 24824–24837. Curran Associates, Inc.
- Yige Xu, Xu Guo, Zhiwei Zeng, and Chunyan Miao. 2025. [Softcot: Soft chain-of-thought for efficient reasoning with llms](#). *Preprint*, arXiv:2502.12134.
- Zeyuan Yang, Xueyang Yu, Delin Chen, Maohao Shen, and Chuang Gan. 2025. [Machine mental imagery: Empower multimodal reasoning with latent visual tokens](#). *Preprint*, arXiv:2506.17218.
- Eric Zelikman, Georges Harik, Yijia Shao, Varuna Jayasiri, Nick Haber, and Noah D. Goodman. 2024. [Quiet-star: Language models can teach themselves to think before speaking](#). *Preprint*, arXiv:2403.09629.
- Huanyu Zhang, Wenshan Wu, Chengzu Li, Ning Shang, Yan Xia, Yangyu Huang, Yifan Zhang, Li Dong, Zhang Zhang, Liang Wang, Tieniu Tan, and Furu Wei. 2025a. [Latent sketchpad: Sketching visual thoughts to elicit multimodal reasoning in mllms](#). *Preprint*, arXiv:2510.24514.
- Juntian Zhang, Chuanqi Cheng, Yuhan Liu, Wei Liu, Jian Luan, and Rui Yan. 2025b. [Weaving context across images: Improving vision-language models through focus-centric visual chains](#). In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 27782–27798, Vienna, Austria. Association for Computational Linguistics.
- Juntian Zhang, Song Jin, Chuanqi Cheng, Yuhan Liu, Yankai Lin, Xun Zhang, Yufei Zhang, Fei Jiang, Guojun Yin, Wei Lin, and Rui Yan. 2025c. [Viper: Empowering the self-evolution of visual perception abilities in vision-language model](#). *Preprint*, arXiv:2510.24285.
- Ziwei Zheng, Michael Yang, Jack Hong, Chenxiao Zhao, Guohai Xu, Le Yang, Chao Shen, and Xing Yu. 2026. [Deepeyes: Incentivizing "thinking with images" via reinforcement learning](#). *Preprint*, arXiv:2505.14362.
- Deyao Zhu, Jun Chen, Xiaoqian Shen, Xiang Li, and Mohamed Elhoseiny. 2023. [Minigt-4: Enhancing vision-language understanding with advanced large language models](#). *Preprint*, arXiv:2304.10592.

A Implementation Details

The model is fine-tuned for 320 steps on 8 MI210 GPUs using 8 gradient accumulation steps. We utilize the AdamW optimizer with $\epsilon = 1e-6$ and a weight decay of 0.1. The learning rate is initialized at $1e-5$ with a cosine decay scheduler and a warmup ratio of 0.03. To optimize memory efficiency, we employ DeepSpeed ZeRO-3 with CPU offloading and enable Flash Attention 2. We adopt a dynamic batching strategy with a token cap of 8,192 and a maximum per-device batch size of 16. For the proposed Laser objective, when the entropy intervention is triggered, the mixing coefficient for the hard target is set to $\alpha = 0.8$. The input image resolution is dynamic, ranging from 128 to 8,192 tokens (approx. 100K to 6.4M pixels).

B Baseline Details

Vision-R1 (Huang et al., 2025). Vision-R1 explores R1-style post-training for MLLMs by constructing a large multimodal chain-of-thought cold start dataset and then applying RL with strategies such as progressive thinking suppression to improve multimodal reasoning.

PAPO (Wang et al., 2026). PAPO (Perception-Aware Policy Optimization) targets the perception bottleneck in multimodal RLVR by introducing an implicit perception loss (KL term) that can be plugged into GRPO/DAPO, improving vision-dependent tasks without extra reward models or teacher models.

DeepEyes (Zheng et al., 2026). DeepEyes studies interleaved multimodal reasoning (“thinking with images”) and incentivizes tool-assisted visual reasoning behaviors via end-to-end RL, using tailored data selection and reward design to improve grounding and reduce hallucinations.

Monet (Wang et al., 2025b). Monet enables latent visual reasoning by generating continuous visual embeddings as intermediate “visual thoughts” and proposes a multi-stage distillation SFT pipeline plus visual-latent policy optimization to better train reasoning in latent visual space.

VL-Rethinker (Wang et al., 2025a). VL-Rethinker enhances slow-thinking for VLMs via RL by adapting GRPO with selective sample replay and enforcing explicit self-reflection steps, improving multi-step multimodal reasoning performance.

LLaVA-OneVision (Li et al., 2024a). LLaVA-OneVision is a unified open multimodal model de-

signed to perform well across single-image, multi-image, and video scenarios, with strong cross-scenario transfer from images to videos.

InternVL3.5-8B (Wang et al., 2025c). InternVL3.5 is a family of open-source multimodal models that improves reasoning and efficiency using cascade RL training and dynamic visual resolution routing; we adopt the 8B variant as a representative strong open baseline.

Qwen2.5-VL-7B (Bai et al., 2025b). Qwen2.5-VL is a flagship open VLM series with improved recognition, localization, and document/video understanding; it introduces dynamic-resolution processing and absolute time encoding for long-video comprehension. We use the 7B model as the base open VLM baseline.

LVR (Li et al., 2025). Latent Visual Reasoning (LVR) enables autoregressive reasoning in the visual embedding space by training the model to generate latent states that reconstruct key visual tokens, interleaved with text generation; it can be further combined with RL to balance latent reasoning and textual outputs.

C Benchmark Details

MMStar (Chen et al., 2024a) is a vision-indispensable multimodal benchmark with carefully curated and purified samples to reduce language-only shortcuts and data contamination. It contains 1.5K questions and evaluates LVLMs over multiple core capabilities with fine-grained axes, providing a balanced diagnosis of both perception and reasoning.

MMVP (Tong et al., 2024) is constructed from 150 CLIP-blind image pairs and probes whether models can truly discriminate basic visual patterns that are obvious to humans but challenging for CLIP-like embeddings. It emphasizes failures where models produce high-confidence yet incorrect answers, often accompanied by plausible-sounding but ungrounded rationales.

BLINK (Fu et al., 2024) is a perception-centric benchmark with 3.8K multiple-choice questions, designed to test visual skills that humans can solve “within a blink,” such as correspondence, depth/geometry cues, forensics, and multi-view reasoning. Its tasks are intentionally difficult to solve from language priors alone, thereby stressing image-dependent perception.

SEEDBench2Plus (Li et al., 2024b) targets text-

rich visual comprehension in real-world formats (Charts, Maps, Web pages). It uses human-annotated multiple-choice questions to assess whether MLLMs can robustly read and reason over dense visual-text content.

Hallusionbench (Guan et al., 2024) is a diagnostic benchmark for entangled language hallucination and visual illusion in LVLMs. It uses human-crafted question structures with a control-group design to quantify logical consistency and common hallucination/illusion failure modes.

HRBench (Wang et al., 2025d) evaluates fine-grained perception on high-resolution (4K/8K) images, where downsampling typically removes crucial details. It systematically tests HR visual understanding (including fine-grained single-/cross-instance perception) to measure whether models can handle true HR content.

D RL Analysis

While our proposed method, Laser, demonstrates robust performance in the supervised fine-tuning stage, we further explore a reinforcement learning framework to align training behavior with inference dynamics and enable an autonomous “early exit” mechanism. To this end, we introduce the **EPG-GRPO** (Expected Policy Gradient - Group Relative Policy Optimization) algorithm. This framework integrates a variance-reduced gradient estimator with a length-invariant policy optimization scheme, successfully guiding the model to significantly reduce token consumption while maintaining performance parity.

D.1 Optimization Objective

Our objective combines Expected Policy Gradients (EPG) to stabilize learning within the high-variance latent space and a modified GRPO formulation to eliminate length bias.

Within LVR regions, standard single-token sampling often leads to gradient instability due to semantic ambiguity. To mitigate this, we calculate the gradient over the expectation of the Top-P subspace (S_{TopP}) rather than a single sampled token. We define the importance ratio ρ_w for each token $w \in S_{\text{TopP}}$ and compute the expected surrogate loss $\mathcal{L}_{\text{epg}}^{(t)}$ as follows:

$$\rho_w = \frac{\pi_\theta(w|s_t)}{\pi_{\text{old}}(w|s_t) + \epsilon}, \quad (9)$$

$$\mathcal{L}_{\text{epg}}^{(t)} = \mathbb{E}_{w \in S_{\text{TopP}}} [\min(\rho_w A_t, \text{clip}(\rho_w, 1 - \epsilon, 1 + \epsilon) A_t)], \quad (10)$$

where A_t is the advantage derived from group sampling.

For global optimization, standard GRPO normalizes by sequence length, which inadvertently incentivizes verbosity. We address this by standardizing the global loss using a fixed maximum completion length L_{max} , independent of the actual trajectory length T_i :

$$\mathcal{L}_{\text{policy}} = \frac{1}{B \times L_{\text{max}}} \sum_{i=1}^B \sum_{t=1}^{T_i} \mathcal{L}_{\text{token}}^{(i,t)} + \beta \cdot \text{KL}(\pi_{\text{ref}} \parallel \pi_\theta). \quad (11)$$

Here, $\mathcal{L}_{\text{token}}$ utilizes the derived \mathcal{L}_{epg} for LVR steps and the standard clipped surrogate loss for explicit tokens.

D.2 Exploration Strategy

We employ a composite exploration strategy to prevent convergence to local optima and robustly learn termination conditions. First, we implement **Relative Norm Perturbation** to enhance the diversity of deterministic LVR hidden states. We inject Gaussian noise scaled by the signal’s norm during the forward pass. For a hidden state \mathbf{h} , the perturbed state \mathbf{h}' is given by:

$$\mathbf{h}' = \mathbf{h} + \lambda \frac{\|\mathbf{h}\|}{\|\epsilon\|} \epsilon, \quad \epsilon \sim \mathcal{N}(0, \mathbf{I}), \quad (12)$$

where we set the noise ratio $\lambda = 0.05$. This ensures the perturbation is significant yet non-destructive across varying signal magnitudes. Second, to encourage efficiency, we apply **Stochastic Horizon Truncation**. For a subset of the generated samples, the maximum allowed steps T_{max} are randomly sampled from a range $[T_{\text{min}}, T_{\text{upper}}]$. This forces the model to attempt convergence within limited horizons, thereby learning to optimize the efficiency of its reasoning path without relying on fixed-length priors.

D.3 Composite Reward Engineering

The reward function R_{total} acts as a multi-objective optimization signal, aggregating components to balance accuracy, structural validity, efficiency, and diversity:

$$R_{\text{total}} = R_{\text{acc}} + R_{\text{fmt}} + R_{\text{eff}} + R_{\text{div}}. \quad (13)$$

Accuracy Reward (R_{acc}). Since our evaluation relies principally on multiple-choice questions or exact answer matching, we assign a binary reward based on the correctness of the final output:

$$R_{acc} = \begin{cases} 1 & \text{if answer is correct,} \\ 0 & \text{otherwise.} \end{cases} \quad (14)$$

Format Reward (R_{fmt}). This component enforces the structural integrity of the dynamic protocol. A reward r_{fmt} is granted solely if the sequence includes the requisite start token, successfully triggers the autonomous termination token $\langle \text{[laser_end]} \rangle$, and correctly encloses the final result in answer tags.

Efficiency Bonus (R_{eff}). To incentivize the model to voluntarily "exit early" when confident, we introduce a dynamic efficiency bonus. This reward is conditional on three strict constraints: the answer must be correct, the trajectory must not be forcibly truncated by the system limit (T_{max}), and the output must be free of format anomalies.

$$R_{eff} = \begin{cases} \beta_{base} - \lambda_{step} \cdot T_{actual} & \text{if Correct} \wedge T_{actual} < T_{max}, \\ 0 & \text{otherwise.} \end{cases} \quad (15)$$

Here, β_{base} represents the maximum potential bonus, and λ_{step} is a penalty coefficient that reduces the reward linearly with each reasoning step used. Under this formulation, a correct answer obtained via system truncation yields zero bonus, explicitly encouraging the model to learn autonomous termination logic.

Diversity Penalty (R_{div}). To prevent "state stagnation"—where adjacent reasoning steps exhibit excessive semantic redundancy—we apply a penalty based on the squared cosine similarity between consecutive hidden states \mathbf{h}_t and \mathbf{h}_{t-1} .

$$R_{div} = -\lambda_{div} \cdot \frac{1}{T-1} \sum_{t=1}^{T-1} (\text{sim}(\mathbf{h}_t, \mathbf{h}_{t-1}))^2. \quad (16)$$

We utilize the squared similarity to impose a non-linear penalty that aggressively punishes high-similarity states while remaining tolerant of the minor correlations necessary for coherent reasoning flow. The term is weighted by λ_{div} and normalized by trajectory length to ensure consistent scaling across varying reasoning depths.

D.4 Experimental Analysis

To validate the efficacy of the EPG-GRPO framework, we compare the supervised baseline (Laser) with its RL-enhanced variant (Laser + EPG). As shown in Table 4, the results demonstrate that our strategy successfully balances computational efficiency with reasoning capability.

The most significant impact is observed in inference efficiency. The RL-enhanced model reduces the average number of generated tokens by approximately 50% across dynamic benchmarks such as BLINK and HRBench. This substantial decrease confirms that the model effectively internalized the autonomous early exit mechanism incentivized by the efficiency bonus R_{eff} and stochastic horizon truncation, learning that concise reasoning paths are often sufficient.

Crucially, this efficiency gain is achieved without compromising general performance. The overall accuracy remains stable, with the method showing particular robustness on hallucinations (HallusionBench) and multi-modal reasoning (MMStar). While minor regressions were observed in select sensitivity-heavy tasks, the Subspace-EPG objective successfully preserved the semantic richness of the latent space, preventing the catastrophic forgetting or mode collapse often associated with RL fine-tuning.

Model	Avg. Tokens ↓		Accuracy (%)						
	BLINK	HRBench	MMStar	MMVP	BLINK	SEED	Hallusion	HRBench	Overall
Laser (Main)	6.03	5.74	60.27	72.00	56.92	70.05	67.72	72.50	66.58
Laser + EPG	3.36	2.87	60.87	72.00	55.76	70.79	68.98	72.00	66.73

Table 4: Comparison between the main model (Laser) and the RL-enhanced model (Laser + EPG). **Avg. Tokens** denotes the average number of generated tokens on BLINK and HRBench.

E Threshold Analysis for an Entropy-Adaptive Mechanism

Threshold (η)	Trigger	MMStar	MMVP	BLINK	SEED	Hallusion	HRBench	Overall
$\eta = 1.0$	0.0%	59.93	69.00	56.86	69.52	67.93	71.88	65.85
$\eta = 0.8$	2.5%	60.33	69.67	56.08	69.96	68.24	71.75	66.01
$\eta = 0.6$ (Ours)	10.0%	60.27	72.00	56.92	70.05	67.72	72.50	66.58
$\eta = 0.5$	18.0%	57.40	71.67	54.97	68.95	64.56	72.75	65.05

Table 5: Ablation study on the entropy threshold η . "Trigger" denotes the intervention activation ratio. $\eta = 0.6$ yields the best balance and is used as our default setting.

E.1 Analysis of Entropy Threshold (η)

We first analyze the impact of the entropy threshold η by correlating it with the **Trigger Ratio**, defined as the percentage of tokens where the model’s high uncertainty necessitates a hard teacher intervention. As shown in Table 5, optimal performance is achieved at $\eta = 0.6$, corresponding to a trigger ratio of approximately 10%. This suggests that intervening on roughly one in ten tokens provides sufficient grounding signals to correct the reasoning trajectory without disrupting the semantic flow.

When the threshold is lowered to $\eta = 0.5$, the trigger ratio rises to 18%, imposing a stricter constraint akin to standard supervision. While this rigid guidance benefits tasks requiring precise alignment (yielding the highest score on HRBench), it stifles the latent exploration necessary for complex logic, resulting in performance degradation on reasoning-heavy benchmarks like MMStar (57.40) and HallusionBench (64.56). Conversely, higher thresholds ($\eta = 0.8, 1.0$) lead to negligible intervention ($< 2.5\%$). Although this preserves flexibility, it lacks the necessary corrective mechanism to fix visual grounding errors, leading to suboptimal results in fine-grained perception tasks such as MMVP. Thus, $\eta = 0.6$ strikes an effective balance, enforcing visual validity while maintaining cognitive flexibility.

Threshold (α)	MMStar	MMVP	BLINK	SEED	Hallusion	HRBench	Overall
$\alpha = 0.2$	60.00	70.00	55.50	69.65	68.66	72.62	66.07
$\alpha = 0.5$	59.87	73.00	55.65	69.83	68.24	72.38	66.50
$\alpha = 0.8$ (Ours)	60.27	72.00	56.92	70.05	67.72	72.50	66.58

Table 6: Ablation study on the parameter α . $\alpha \in [0, 1]$ controls the intensity of the hard intervention. We observe that $\alpha = 0.8$ achieves the optimal balance.

E.2 Analysis of Intervention Intensity (α)

We further examine the impact of the parameter α , which modulates the intensity of the hard intervention once high uncertainty is detected. As indicated in Table 6, increasing α leads to a consistent, albeit marginal, improvement in overall performance. We observe that a lower α (e.g., 0.2) applies a softer correction, which appears insufficient to fully resolve ambiguity when the model is highly uncertain. Conversely, a higher setting ($\alpha = 0.8$) provides a more decisive guidance signal, effectively acting as a “hard reset” to realign the latent trajectory with the ground truth. This suggests that during critical states of high uncertainty, prioritizing deterministic

constraints over soft superposition yields slightly more robust reasoning.

F Dataset Details

Source	Task Domain	Data Scale		Reasoning Nodes		
		Count	Ratio	Min	Max	Mean
Flickr30k	Captioning	103,790	38.5%	2	20	7.50
GQA	VQA/Reasoning	86,218	32.0%	2	18	7.13
OpenImages	Detection	42,639	15.8%	3	17	6.38
Visual7W	VQA	30,271	11.2%	1	17	6.85
CUB	Fine-grained Cls.	3,521	1.3%	4	18	8.55
VSR	Spatial Reasoning	3,334	1.2%	2	16	7.08
Total / Avg.	-	269,773	100%	1	20	7.14

Table 7: Detailed statistics of the *ScanPath* dataset. The dataset integrates six diverse visual tasks. **Nodes** refers to the number of discrete semantic anchors in the reasoning chain (distinct from token length).

Table 7 outlines the detailed statistics of the *ScanPath* dataset, comprising 269,773 samples across six visual domains. We measure chain complexity using “Reasoning Nodes”—defined as discrete semantic anchors (e.g., region identification or attribute verification) rather than linguistic token length. As shown in the table, the node distribution naturally aligns with task difficulty: fine-grained tasks like CUB-200 require deeper reasoning paths (Mean: 8.55 nodes) compared to basic detection tasks like OpenImages (Mean: 6.38 nodes), resulting in a balanced overall average of 7.14 nodes.

G Time-Aware Semantic Decay

We further explore a **Time-Aware Semantic Decay** component to regulate the semantic distribution within the validity window. This strategy applies a temporal bias to the target logits based on their relative distance from the current step:

$$\tilde{z}_t^{(k)} = \hat{z}_t^{(k)} + \ln(\gamma^{k-t}), \quad (17)$$

where $k \in W_t$ denotes the token index and $\gamma \in (0, 1]$ acts as a decay factor. This formulation allows the framework to flexibly modulate the attention density assigned to distant future tokens relative to immediate deductive steps.

Model	MMStar	MMVP	BLINK	SEED	Hallusion	HRBench	Overall
Larse (Default)	60.27	72.00	56.92	70.05	67.72	72.50	66.58
w/ Time Decay	59.60	73.00	55.34	70.18	67.40	72.38	66.32

Table 8: Ablation study comparing the standard Laser model with a time-aware variant.

H Prompt Engineering

To ensure the synthesis of high-quality cognitive scanpaths, we employed the following structured system prompt. It enforces a strict “Global-to-Local” scanning logic and output format constraints.

```
### Role Definition
You are a Visual Cognitive Engine designed to deconstruct
the visual reasoning process into a strictly ordered "Visual
Scanpath".

### Core Objective
Generate a sequential stream of Atomic Visual Concepts.
This stream must represent a logical flow of discovery:
scanning from the global environment, zooming into specific
objects, and accumulating visual cues, culminating in the
most critical information needed to answer the user's query.

### Precision & Format Principles (CRITICAL)
1. Atomic Specificity: Be as specific as the image clarity
allows immediately (e.g., "Ferrari" not "Car"), but
strictly use single words or 1-3 word phrases.
2. De-Grammatized Output: Output dense information only.
REMOVE all stop words (is, the, a, of, in).
3. Visual Certainty: Only output concepts that are visually
observable. Use broader terms if blurry.
4. Contextual Anchoring: Repeat a previous entity ONLY if
necessary to attach the final resolution.

### The 4-Stage Scanning Logic (Strict Order)
1. Global Anchor (Step 1-2): Start with the broadest visible
context (e.g., "Kitchen", "Blue Sky").
2. Subject Localization (Step 3-X): Locate the main subject
relevant to the question.
3. Visual Evidence (Step X-Y): List visible attributes or
actions supporting the answer.
4. Critical Resolution (Final Step): The specific concept
answering the query must appear at the very end.

### Negative Constraints
- NO Premature Reveals: Do not output the answer early.
- NO Artificial Hierarchy: No "Fruit" → "Apple", just
"Apple".
- NO Sentences: Raw concepts only.

### Output Format
Output strictly valid JSON:
{
  "reasoning_chain": [
    "String1",
    "String2",
    ...
  ]
}
```

clearly demonstrate the efficiency and accuracy of our method during the reasoning process.

I Details of Human Annotations

To assess the quality of our dataset, we recruited three expert annotators (Ph.D. candidates in Computer Science) to manually evaluate 200 randomly sampled instances from *ScanPath*. The evaluation focused on two key dimensions: (1) the validity of the visual reasoning chain, and (2) adherence to the global-to-local logic. The results demonstrated a validity rate of 91.5% with substantial inter-annotator agreement (Fleiss' $\kappa = 0.677$), confirming the reliability of our automated pipeline and filtering protocol.

J Case Study

We present test cases of our method during inference, as illustrated in Figures 6 and 7. These cases

Question: What is the approximate time of day in the image?

Options:


- A: Morning,
- B: Evening,
- C: Noon,
- D: Night,

Answer: "C",



Model_Output: "<|laser_start|> Sun Forest Morning Shadows Noon <|laser_end|><answer> C </answer>"

Figure 6: A test case from MMStar showcases the efficacy and efficiency of our Laser.



Question: Which image shares the same style as the first image?

Option:

- A. the second image
- B. the third image

Answer: "(A)",

Model_output: "<|laser_start|> Street Art Reference Painting Abstract Third Cartoon Second<|laser_end|><answer> A </answer>"

Figure 7: This multi-image reasoning test case from MMStar illustrates the effectiveness and efficiency of our Laser.